

10-12-00

A

BER-016

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of ) Art Unit: unknown  
ARTES )  
Serial No. unknown ) Examiner: unknown  
Filed: herewith )  
Docket No: BER-016  
For: A METHOD FOR TRANSMISSION OF DATA

JCS92 U.S. PRO  
09/666273  
10/10/00

Box Patent Application  
Honorable Commissioner of  
Patents and Trademarks  
Washington, D.C. 20231

TRANSMITTAL LETTER

Sir:

Enclosed for filing please find the patent application for an invention entitled A METHOD FOR TRANSMISSION OF DATA, filed on behalf of JARI ARTES including 11 pages of specification, which includes a one page abstract and 10 total claims of which there are 2 independent claims, and 8 dependent claims, and 5 sheets of Formal Drawings labeled on reverse side comprising Figures 1-5, an unsigned Declaration for Patent Application, and a certified copy of the priority document Finish patent 19992188. The attorney's Docket Number is BER-016.

Kindly address all communications regarding this application to:  
Ronald Craig Fish  
FALK & FISH  
P.O. Box 2258  
Morgan Hill, California 95038  
TEL 408 778-3624  
FAX 408 776-0426

All filing fees and surcharge will be paid when responding to Notice to File Missing Parts of Application.

The Commissioner is authorized to charge any shortfall in fees and any fees necessary at any time to keep this case from becoming abandoned to deposit account 06-0932, or to credit any overpayments to deposit account 06-0932.

A return receipt postcard is enclosed herewith.

Respectfully submitted,

Dated: October 10, 2000

Ronald C. Fish

Ronald Craig Fish  
Reg. No. 28,843

I hereby certify that all of the above identified correspondence  
is being deposited with the United States Postal Service as  
Express Mail in an envelope addressed to: COMMISSIONER OF  
PATENTS AND TRADEMARKS, WASHINGTON, D.C. 20231 on  
October 10, 2000.

Express Mail Receipt Number: EM429087551US

Date: October 10, 2000

Ronald C. Fish

Ronald Craig Fish, President  
Reg. No. 28,843  
Ronald Craig Fish, a Law Corporation

09686273 - 101000

## A method for transmission of data

## BACKGROUND OF THE INVENTION

5

### 1. Field of the Invention

The invention relates to methods for transmission of data, more particularly for transmission of data in clustered structures in IP networks. Especially, the invention is related to such a method as specified in the preamble of the independent method claim.

## 2. Description of Related Art

15 A gateway is typically understood as a device, which connects a first network to a second network. A server is typically understood as a device, which receives requests from a client, performs the requested act and sends results back to the client. Typically, a gateway is used to connect the local area network (LAN) of a company to the public Internet. The gateway typically acts as a firewall, i.e. screens  
20 incoming traffic according to a certain set of rules to protect the local area network and data stored therein. Servers provide services to individual computers within the local area network, such as database services, disk services, and so on.

25 Gateway clusters and server clusters are used, when the capacity of a single device  
is not sufficient. Compared to a single device, gateway or server clusters provide  
high availability and increased throughput. Figure 1 illustrates the structure of a  
gateway cluster. Figure 1 shows client computers 10 connected to a switch 20. The  
client computers 10, the switch 20, and form a local area network (LAN). The LAN  
is connected to an external network 50 through a gateway cluster consisting in this  
30 example of five gateways 30. The gateways 30 may for example function as firewall  
devices. The gateways 30 are connected to the external network 50 through a switch  
20 and a router 40. The client computers are in figure 1 divided into three  
subnetworks A, B, and C.

35 Figure 2 illustrates a server cluster. Figure 2 shows client computers 10 connected to a switch 20, which together form a local area network (LAN). The switch connects the local area network to the server cluster, which in this example comprises five server computers 30. Figure 2 further illustrates mass storage units 60 connected to

the server computers 30. The client computers are in figure 2 divided into three subnetworks A, B, and C.

In general, there are two basic methods to construct server and gateway clusters in a TCP/IP network. According to the first basic method, each server or gateway has a distinct IP address, and an external device redirects the requests or packets to different servers or gateways. The external device has a certain set of rules for controlling the redirection of data packets. For example, the external device can redirect data packets on the basis of current load on the cluster devices. In the example of figures 1 and 2, the external device can be the switch 20 connecting the LAN to the cluster. The main drawback of this solution is asymmetric routing: because routing in IP networks is usually based only on the destination IP address, outgoing and incoming packets can be routed to different gateways. This is a problem when the gateways also serve as firewalls, since firewalls typically do not accept any incoming messages, which are not a reply to a previous outgoing request. When the incoming packets are routed to a second gateway which is a different gateway than the first gateway which sent the original outgoing request, they will be discarded since the second gateway does not have knowledge about the previous request.

According to the second basic method, all units of the cluster share a common IP address and MAC (media access control) address, and therefore receive all traffic in both directions. The units of the cluster have filter functions at their inputs and outputs, which according to certain rules allow the processing of a packet by only one of the gateways, while the others ignore the packet. The rules are constructed in such a way that the same gateway processes the packets in both directions of a single connection, which allows normal firewall operation. The benefit of this solution is transparency and scalability. New servers or gateways can be added to the cluster without changing routing, i.e. unnoticed by the clients.

However, even the second basic method has its drawbacks. Connecting multiple servers or gateways sharing a common MAC address to a switch requires usually usage of a multicast MAC address for the cluster, i.e. a hardware layer address which is recognized and processed by the network interface of every unit of the cluster. The network interfaces of the gateways always have individual MAC addresses as well, but in this solution, the interfaces are arranged to recognize the multicast MAC address selected for the cluster. The IP address of the server or gateway cluster is mapped to a multicast MAC address in the ARP (Address

Resolution Protocol) tables of the nodes of the IP network. When a client wishes to send a data packet to the cluster, it addresses the packet to the IP address of the cluster. The network translates the IP address to a MAC address for transportation on the physical layer, and in this case the MAC address is a multicast address. The

- 5 problem with multicast MAC addresses is that switches typically flood packets having a multicast address as destination to all ports of the switch by default, wherefore the packet ends up in all other subnetworks connected to other ports of the switch. This effect causes an unnecessary loading of other subnetworks connected to other ports of the switch than the gateways/servers of the cluster. This  
10 is a severe problem, if the LAN is divided into virtual LAN networks (VLAN).

Virtual LANs (VLAN) can be created using switches to segment networks. A single segment is a broadcast domain, i.e. broadcast, multicast and unknown unicast frames are sent to all stations in the segment. Port-based virtual LAN networks are groups of switch ports and attached segments (subnetworks) which belong to the VLAN.

#### SUMMARY OF THE INVENTION

- 20 An object of the invention is to realize a method for transmission of data in an IP packet network, which avoids the problems of the prior art.

The objects are reached by using the IGMP protocol to inform the switch, which of its ports are connected to units of a cluster system, and forwarding those packets  
25 with the unicast IP address of the cluster system and the multicast MAC address of the cluster system only to those ports.

The method according to the invention is characterized by that, which is specified in the independent method claim. The cluster system according to the invention is  
30 characterized by that, which is specified in the independent claim directed to a cluster system. The dependent claims describe further advantageous embodiments of the invention.

The invention is applicable to gateway clusters as well as server clusters. In the  
35 following, the single unit of cluster i.e. a gateway or a server is denoted a cluster unit.

According to the invention, the cluster units are configured to be members of an IP multicast group specific to the cluster. The switch or switches directly connected to the cluster units are arranged to monitor multicast group membership reports from the cluster units, and therefore obtain knowledge about which ports of the switch or

- 5 switches are connected to cluster units. Advantageously, the switch or switches may also send membership queries to find out, which ports are connected to members of the cluster multicast group. Consequently, when the switch receives a packet with a multicast MAC address and the unicast IP address of the cluster, the switch sends the packet to only those ports to which cluster units are connected, and not to all  
10 ports of the switch as according to the prior art. In the local area network, the unicast IP address of the cluster is mapped to a multicast MAC address in the routing tables in the network. Therefore, when any client sends a packet to the cluster, the switch will receive the packet with a multicast MAC destination address.

- 15 For obtaining this functionality, the switch is arranged to recognize the unicast IP address of the cluster and the multicast group address configured for the cluster, and to store into its memory information about which ports are connected to such network nodes, which send group membership reports for the multicast group address configured for the cluster.

- 20 According to the IGMP protocol (Internet Group Management Protocol), network nodes send a group membership report when they join a multicast group to announce their presence to multicast routers. The report is sent with an IP destination address equal to the host group address being reported, so that the  
25 multicast router can recognize which multicast groups have members in its domain. A multicast router sends an IGMP query at regular intervals to see if any nodes belong to any multicast groups. Network nodes send a group membership report, when they receive a group membership query from the router. According to the invention, this mechanism is used to inform the switch about which ports are  
30 connected to cluster units. The IGMP protocol version 1 is defined in the IETF standard document RFC 1112, and IGMP version 2 is defined in the IETF proposed standard document RFC 2236.

#### BRIEF DESCRIPTION OF THE DRAWINGS

35

The invention is described in more detail in the following with reference to the accompanying drawings, of which

Figure 1 illustrates a gateway cluster arrangement according to prior art,

Figure 2 illustrates a server cluster arrangement according to prior art.

5 Figure 3 illustrates a method according to an advantageous embodiment of the invention.

Figure 4 illustrates another method according to an advantageous embodiment of the invention, and

19

Figure 5 illustrates a cluster system according to an advantageous embodiment of the invention.

Same reference numerals are used for similar entities in the figures.

15

## DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

#### A. A FIRST GROUP OF ADVANTAGEOUS EMBODIMENTS

20 Figure 3 illustrates a method according to an advantageous embodiment of the invention. Figure 3 illustrates a method for transmission of data in an IP packet network, which network comprises a cluster of cluster units and a switching unit having a plurality of ports, and in which network the cluster units are connected to a part of the plurality of ports, and which cluster units share a unicast IP address.

25

According to the present embodiment, the method comprises at least steps, in which

- the cluster units are configured 100 to be members of an IP multicast group specific to the cluster.

- the IGMP protocol is used to obtain 110 information about to which ports of the plurality of ports the cluster units are connected.

- the MAC address of a received IP packet is checked in step 140, and if said MAC address is found to be a multicast MAC address, the IP destination address of said packet is compared 150 to the unicast IP address shared by the cluster units,

35 - if the IP destination address of said packet is the same as the unicast IP address, the packet is forwarded 160 to those ports, to which the cluster units were found to be connected, and if it is not, the packet is sent to all ports of the switch.

Figure 4 illustrates as an example how the step 110 of obtaining information is performed according according to an advantageous embodiment of the invention. According to this embodiment, the method further comprises at least steps, in which the switching unit receives 114 an IGMP group membership report, the switching unit checks 116, if said report is addressed to said IP multicast group specific to the cluster, and if it is, the switching unit stores 118 into its memory the identifier of the port, through which said report arrived.

Advantageously, the method may as well comprise a step, in which an IGMP group membership query is sent 112 by the switching unit.

In the inventive method, the IGMP protocol version 1 according to RFC 1112 can be used, as well as the IGMP protocol version 2 according to RFC 2236.

## 15 B. A SECOND GROUP OF ADVANTAGEOUS EMBODIMENTS

Figure 5 illustrates a cluster system according to an advantageous embodiment of the invention. Figure 5 shows client computers 10 connected to a switch 200. The client computers 10, the switch 200, and form a local area network (LAN). The LAN is connected to an external network 50 through a gateway cluster consisting in this example of five gateways 30. The gateways 30 may for example function as firewall devices. The gateways 30 are connected to the external network 50 through a switch 200 and a router 40. The client computers are in figure 5 divided into three subnetworks A, B, and C. According to the embodiment shown in figure 5, the cluster system has a plurality of cluster units 30 and a switching unit 200, and said cluster units are associated with the same IP unicast address. According to the embodiment, the cluster units are configured to be members of an IP multicast group specific to the cluster system, and the cluster system comprises

- means 240 for observing using the IGMP protocol which ports of the switching unit are connected to the cluster units,

- means 210 in the switching unit for observing the MAC destination address of a packet arriving to the switching unit and for checking if said MAC destination address is a MAC multicast address,

- means 220 in the switching unit for observing the IP destination address of said packet and for comparing said IP destination address to said IP unicast address associated with the cluster units,

- means 230 in the switching unit for forwarding of the packet to those ports whose identifiers were previously stored to said memory means as a response to the finding

that said IP destination address and said IP unicast address are the same and said MAC destination address is a MAC multicast address.

In a further advantageous embodiment of the invention, said means 240 for observing using the IGMP protocol comprise at least

- means 242 in the switching unit for observing IGMP multicast group reports and for checking, if a received IGMP multicast group report is addressed to said IP multicast group specific to the cluster system, and

- means 244 in the switching unit for storing into a memory means an identifier of that port via which said received IGMP multicast group report arrived as a response to finding that said report was addressed to said IP multicast group.

In a still further advantageous embodiment of the invention, the system further comprises means 260 in the switching unit for sending IGMP group membership queries.

In an advantageous embodiment of the invention, said means 210, 220 ,230, 240, 242, 244, 250, and 260 are realized using software programs stored in a memory element 270 of the switching unit 200 and executed by a processor 280 of the switching unit 200.

In another advantageous embodiment of the invention, said means 210, 220 ,230, 240, 242, 244, 250, and 260 are realized using dedicated logic circuits implemented for example using an ASIC circuit (application specific integrated circuit). Such an implementation of the invention allows the functioning of the switching unit at extremely high speeds.

In an advantageous embodiment of the invention, the cluster units 30 are gateway units 30.

In a further advantageous embodiment of the invention, the cluster units are server units.

#### C. FURTHER CONSIDERATIONS

The invention has several advantages. For example, the invention enables connecting gateway clusters to switches, which allows the construction of high-speed networks using port-based virtual local area networks.

The invention is applicable with IPv4 as well as IPv6 networks. The IP protocol version 4 is described in detail in the specification RFC 791. The next version of the IP protocol, known as IPv6, is described in the specification RFC 1883.

5

In view of the foregoing description it will be evident to a person skilled in the art that various modifications may be made within the scope of the invention. While a preferred embodiment of the invention has been described in detail, it should be apparent that many modifications and variations thereto are possible, all of which fall within the true spirit and scope of the invention.

Claims

✓. Method for transmission of data in an IP packet network, which comprises a cluster of cluster units, a switching unit having a plurality of ports, the cluster units

5 being connected to a part of the plurality of ports, which cluster units share a unicast IP address, said method comprising at least steps, in which

- the cluster units are configured to be members of an IP multicast group specific to the cluster,

10 - the IGMP protocol is used to obtain information about to which ports of the plurality of ports the cluster units are connected,

- the MAC address of a received IP packet is checked, and if said MAC address is a multicast MAC address, the IP destination address of said packet is compared to the unicast IP address shared by the cluster units,

15 - if the IP destination address of said packet is the same as the unicast IP address, the packet is forwarded to those ports, to which the cluster units were found to be connected.

2. The method of claim 1, further comprising at least steps, in which the switching unit receives an IGMP group membership report,

20 the switching unit checks, if said report is addressed to said IP multicast group specific to the cluster,

and if it is, the switching unit stores into its memory the identifier of the port, through which said report arrived.

25 3. The method of claim 1, further comprising a step, in which an IGMP group membership query is sent by the switching unit.

4. The method of claim 1, wherein the IGMP protocol is protocol version 1 according to RFC 1112.

30 5. The method of claim 1, wherein the IGMP protocol is protocol version 2 according to RFC 2236.

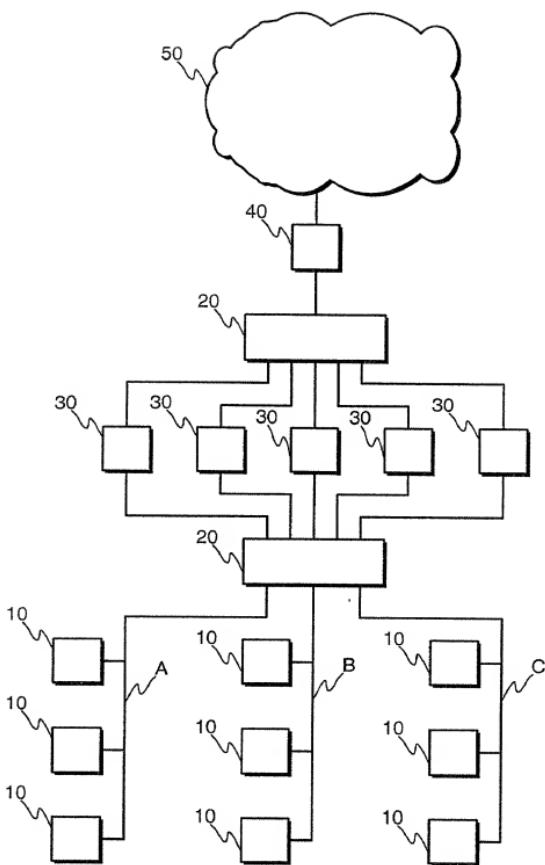
6. Cluster system having a plurality of cluster units and a switching unit, said cluster units being associated with the same IP unicast address, wherein the cluster units are configured to be members of an IP multicast group specific to the cluster system, said cluster system comprising
- 5 - means for observing using the IGMP protocol which ports of the switching unit are connected to the cluster units,
- means in the switching unit for observing the MAC destination address of a packet arriving to the switching unit and for checking if said MAC destination address is a MAC multicast address,
- 10 - means in the switching unit for observing the IP destination address of said packet and for comparing said IP destination address to said IP unicast address associated with the cluster units,
- means in the switching unit for forwarding of the packet to those ports whose identifiers were previously stored to said memory means as a response to the finding that said IP destination address and said IP unicast address are the same and said MAC destination address is a MAC multicast address.
- 15
7. The system according to claim 6, wherein said means for observing using the IGMP protocol comprise at least
- 20 - means in the switching unit for observing IGMP multicast group reports and for checking, if a received IGMP multicast group report is addressed to said IP multicast group specific to the cluster system, and
- means in the switching unit for storing into a memory means an identifier of that port via which said received IGMP multicast group report arrived as a response to finding that said report was addressed to said IP multicast group.
- 25
8. The system according to claim 6, further comprising means in the switching unit for sending IGMP group membership queries.
- 30 9. The system according to claim 6, wherein the cluster units are gateway units.
10. The system according to claim 6, wherein the cluster units are server units.

DOCUMENT NUMBER

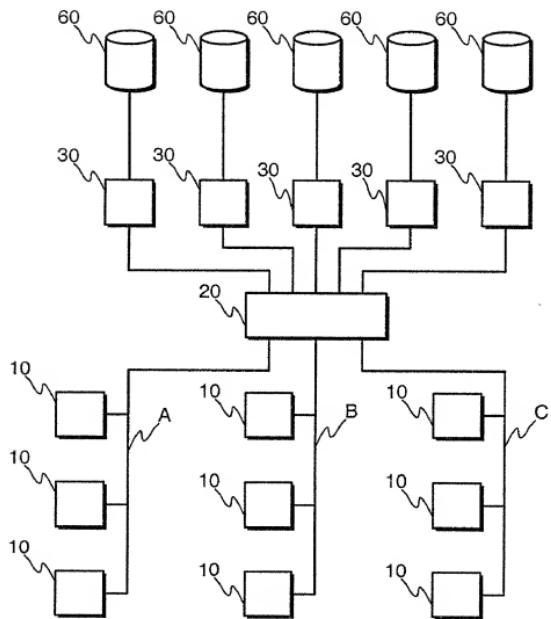
## ABSTRACT

The invention relates to methods for transmission of data, more particularly for transmission of data in clustered structures in IP networks. According to the invention, the cluster units are configured to be members of an IP multicast group specific to the cluster. The switch or switches directly connected to the cluster units are arranged to monitor multicast group membership reports from the cluster units, and therefore obtain knowledge about which ports of the switch or switches are connected to cluster units. Advantageously, the switch or switches may also send membership queries to find out, which ports are connected to members of the cluster multicast group. Consequently, when the switch receives a packet with a multicast MAC address and the IP address of the cluster, the switch sends the packet to only those ports to which cluster units are connected, and not to all ports of the switch as according to the prior art.

Figure 3



**Fig. 1**  
PRIOR ART



**Fig. 2**

PRIOR ART

00000000000000000000000000000000

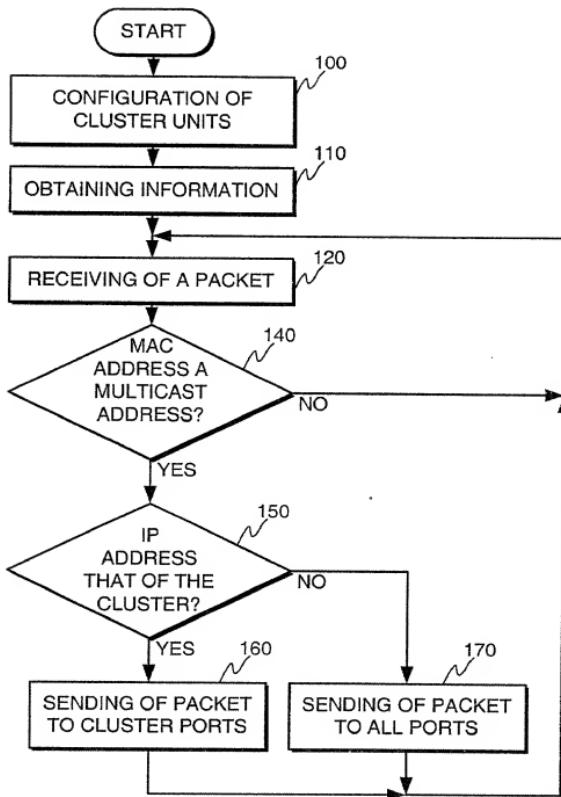
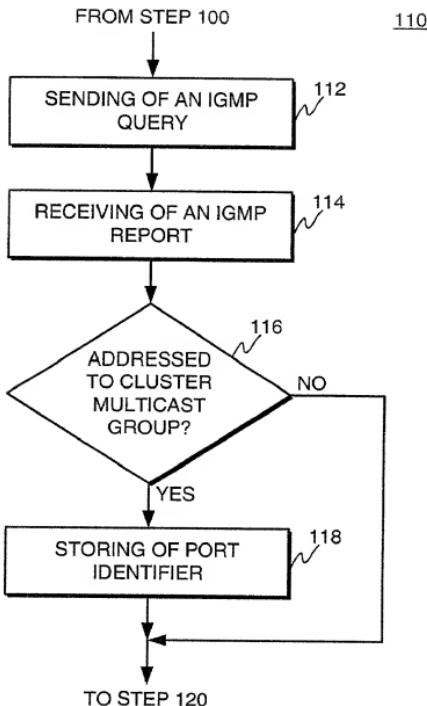
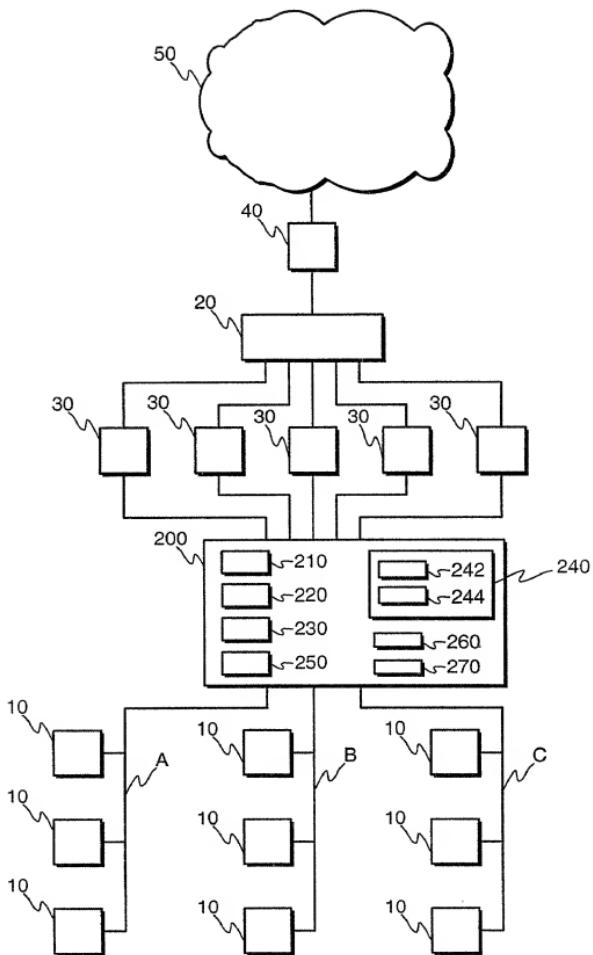


Fig. 3



**Fig. 4**



**Fig. 5**

**DECLARATION FOR PATENT APPLICATION**

Inventor(s): JARI ARTES  
Docket No. BER-016

As a below-named inventor, I (We) hereby declare that:

My correct residence, post office address and citizenship are stated below next to my name.

I believe myself to be the original, first and sole inventor (if only one name is listed below) or an original and first joint inventor (if more than one name is listed below) of the subject matter which is disclosed and claimed and for which a patent is sought on the invention entitled:

**A METHOD FOR TRANSMISSION OF DATA**

The specification of this subject matter:

- is attached hereto.
- was filed on \_\_\_\_\_
- was assigned serial No. \_\_\_\_\_
- which was amended on \_\_\_\_\_

I hereby state that I have reviewed and understand the contents of the above identified patent application, including the claims, as amended by any amendment(s) referred to above. I believe the subject matter claimed in the above-identified application to be new and to be unobvious to persons of ordinary skill in the art in view of the prior art of which I am aware. I further hereby state that the specification of the above identified patent application adequately describes how to make and use the claimed invention, and further that it sets forth the best mode for practicing the invention known to me as of the date that the application was filed. I acknowledge the duty to disclose information which is material to the examination of this application in accordance with 37 C.F.R. 1.56. I further acknowledge the duty to disclose information known to me to be material under 37 C.F.R. 1.56 to the examination of the claims in any continuation-in-part application filed under the conditions specified in 35 U.S.C. Section 120 which became available between the filing date of the prior application or the national or PCT international filing date and the filing date of the continuation-in-part application.

I hereby claim foreign priority benefits under 35 U.S.C. §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed.

Application No.	Country	Filing Date	Priority Claimed
19992188	Finland	11 October 1999	yes

I hereby claim the benefit under 35 U.S.C. §120 of any United States application(s) listed

below and, insofar as the subject matter of each of the claims of this application is not disclosed in these prior United States application(s) in the manner provided by 35 U.S.C. §112, I acknowledge the duty to disclose material information as defined in 37 C.F.R. §1.56(a) which occurred between the filing date of the prior application(s) and the national or PCT international filing date of this application.

Application No.	Filing Date	Status (Issued, Pending, Abandoned)		
FULL NAME OF INVENTOR 1	FIRST Name JARI	MIDDLE Initial(s)	LAST Name ARTES	
RESIDENCE AND CITIZENSHIP	City HELSINKI	State or Foreign Country FINLAND	Country of Citizenship FINLAND	
POST OFFICE ADDRESS	Number and Street ABRAHAMINKATU 17 C 50	City FIN-00180 HELSINKI	State or Country FINLAND	Zip Code

I hereby appoint the following attorney to prosecute this application and to transact all business in the United States Patent and Trademark Office connected with this application and all continuations, divisional and continuations-in-part of this application and to transact all business in all foreign patent offices for all foreign counterparts of this patent application or any divisional, continuation, continuation-in-part application or any PCT application arising out of this patent application or claiming priority of this patent application:

Ronald C. Fish (Reg. Number 28,843)

Address all telephone calls to Ronald C. Fish at telephone no. (408) 778-3624

Address all correspondence to Ronald Craig Fish

FALK & FISH  
P.O. Box 2258  
Morgan Hill, California 95038, U.S.A.  
FAX No. (408) 776-0426

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Title 18, United States Code, § 1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

---

Signature of Inventor 1

Date